

MULTI-PROCESSOR SYSTEM AND METHOD
OF ACCESSING DATA THEREIN

CROSS REFERENCE TO RELATED APPLICATIONS

This application is a continuation of United States Application Serial No. 09/418,520 filed October 15, 1999.

5 TECHNICAL FIELD OF THE INVENTION

The present invention is related in general to computing system architectures and more particularly to a multi-processor system and method of accessing data therein.

BACKGROUND OF THE INVENTION

Controlling access to memory in a multi-processor system is a difficult process, especially when many processors share data in memory. Typically, each processor maintains a small cache of most frequently used data for quick access so that time consuming requests for data to the common system memory may be avoided. However, the cache for each processor must be updated with changes made to its associated data that are reflected in the common system memory. One technique for updating processor caches is to couple each processor to what's known as a snoopy bus. A request for access to data by a requesting processor is broadcast to other processors over the snoopy bus. Each processor "snoops" into their cache to see if it has the most recent copy of the requested data. If a processor does have a most recent copy of the requested data, then that processor provides the data to the requesting processor. If no processor has a most recent copy of the requested data, a memory access is required to fulfill the requesting processor's request. If a processor updates a memory location, this update is broadcasted over the snoopy bus to the other processors in the system. Each processor checks its cache to see if it has the data corresponding to the updated memory location. If so, the processor may either remove that data and corresponding memory location from its cache or update its cache with the new information. This snoopy bus technique is effective for a small number of processors within a computer system but is ineffective for computer systems having hundreds of processors.

Another technique is to provide a directory based memory configuration. For directory based memories, a directory is used to maintain a directory entry corresponding to every entry in memory. The directory entry specifies whether the associated data in memory is valid or where the most recent copy of the data may be accessed. The directory based memory configuration avoids coupling all the processors in the computer system together and having processors be bothered handling broadcast requests found in snoopy bus designs. Communication only needs to occur with the processor having the most recent copy of the data. The size of the directory provides the constraint for this configuration, as the directory would become too large to support the number of processors and memories in a large computer system. Therefore, it is desirable to provide a memory access control mechanism for computer systems with a large number of processors.

SUMMARY OF THE INVENTION

From the foregoing, it may be appreciated that a need has arisen for providing a multi-processor system with processors having integrated memories and memory directories linked together through an external directory. In accordance with the present invention, a multi-processor system and method of accessing data therein are provided that substantially eliminate or reduce disadvantages and problems of conventional multi-processor systems.

According to an embodiment of the present invention, there is provided a multi-processor system that includes a plurality of processors, wherein each processor includes an integrated memory, an integrated memory controller, and an integrated memory directory. The integrated memory provides, receives, and stores data. The integrated memory controller controls access to and from the integrated memory. The integrated memory directory maintains a plurality of memory references to data within the integrated memory. The multi-processor system also includes an external switch coupled to each of the plurality of processors. The external switch passes data to and from any of the plurality of processors. The external switch includes an external directory. The external directory provides a memory reference to remote data for each of the plurality of processors that is not provided within its own integrated memory directory.

The present invention provides various technical advantages over conventional multi-processor systems. For example, one technical advantage is to integrate memory, memory control, and memory directory into a

processor. Another technical advantage is the ability to extend the integrated memory directory capability with external support in order to implement large cache coherent multi-processor systems. Yet another technical advantage is to remove large system directory policy decisions from the individual processor in the system. Still another technical advantage is to provide a directory protocol that can be used with commodity processors having integrated memories and directories. Other technical advantages may be readily ascertainable by those skilled in the art from the following figures, description, and claims.

BRIEF DESCRIPTION OF THE DRAWINGS

For a more complete understanding of the present invention and the advantages thereof, reference is now made to the following description taken in conjunction with the accompanying drawings, wherein like reference numerals represent like parts, in which:

FIGURE 1 illustrates a block diagram of a multi-processor system;

FIGURE 2 illustrates a block diagram of a processor within the multi-processor system; and

FIGURE 3 illustrates a block diagram of an alternate embodiment of the multi-processor system.

DETAILED DESCRIPTION OF THE INVENTION

FIGURE 1 is a block diagram of a multi-processor system 10. Multi-processor system 10 includes a plurality of processors 12 and an external switch 14. Each of the plurality of processors 12 has a memory 16, a memory directory 18, and a central processing unit 20 all integrated into a single device. External switch 14 includes an external directory 22. Each processor 12 may couple to external switch 22 in order to exchange among each other data stored in their respective memories. External switch 22 may also couple to another external switch 22 in order to enlarge the capabilities of multi-processor system 10.

In operation, memory directory 18 of a particular processor 12 includes memory references to data stored within its corresponding memory 16. For smaller multi-processor systems, memory directory 18 may also include memory references to data stored in a remote memory 16 associated with a different processor 12 within a local regional group. As memory sizes and systems become larger, an individual memory directory 18 of a particular processor 12 may not be able to include a memory reference to all data in the system which the particular processor 12 desires to access. In order to alleviate this situation, external directory 22 of external switch 14 includes a capability to retrieve memory references to data in memories remote from the particular processor 12.

When the particular processor 12 desires to access data from a remote memory 16, its memory directory 18 determines that it does not have a memory reference to the desired data. Memory directory 18 generates a data request that is sent to external directory 22 in external

switch 14. External directory 22 processes the request and generates a memory reference to the desired data. External switch 14 uses the generated memory reference to retrieve the desired data and provide it to the requesting processor 12.

Memory directory 18 preferably holds memory references to data that has been most recently accessed. If data is requested by the particular processor 12 and that data resides in its associated memory 16, then memory directory 18 generates a memory reference to the new data. If memory directory 18 is fully occupied with memory references, then memory directory 18 may overwrite the memory reference to data that has not been accessed for the longest period of time with the newly generated memory reference. External directory 22 may operate in a similar manner by maintaining memory references to most recently accessed data from among the plurality of processors 12 and only generate a new memory reference for a request to data not currently represented by a memory reference within external directory 22. Though not necessary, memory references within each memory directory 18 may be represented in a similar manner as memory references in external directory 22.

FIGURE 2 is a block diagram of a processor 12. Processor 12 includes memory 16, a memory controller 30, memory directory 18, one or more network interfaces 32, and a CPU controller 34. Network interfaces 32 provide a communication capability between processor 12 and external switch 22. Memory controller 30 controls the read and write access from and to memory 16. CPU controller 34 controls flow between one or more processing units.

The size of memory directory 18 may vary according to the size of its associated memory 16. For example, a processor 12 holding eight megabytes with sixty-four byte lines of cache in a four to one ratio may use 2(17) entries. Using a four gigabyte dynamic random access memory for memory 16, memory references may be represented by thirteen bit tags, two state bits, four pointer/vector bits and two error correction code (ECC) bits. With twenty-one bits per entry and 2(17) entries, memory directory 18 has a size of 2.6 Megabytes. As another example, a processor 12 holding thirty-two megabytes with one hundred twenty-eight byte lines of cache in a four to one ratio may use 2(18) entries. Using an eight gigabyte dynamic random access memory for memory 16, memory references may be represented by twelve bit tags, two state bits, four pointer/vector bits and two ECC bits. With twenty bits per entry and 2(18) entries, memory directory 18 has a size of 5 Megabytes.

With the presence of external directory 22, each memory directory 18 may be set up to track its local memory 16 cached memory references. External directory 22 may be set up to track remote cached memory references for the processors 12. Through the use of memory directories 18 and at each processor 12 and external directories 22 in a large multi-processor system 10 environment, cache coherency is provided to ensure that all processors 12 have an accurate view of the entire system memory. Requests for memory may even be passed from one external switch 14 to another to further extend the memory and access mechanism of multi-processor system 10.

FIGURE 3 shows an alternate embodiment of multi-processor system 10. In this embodiment, processors 12 are coupled to two external switches 14. The two external switches 14 provide two routing planes for memory access and coherence. The two routing planes may provide redundancy for multi-processor system 10 or extend the bandwidth capability of multi-processor system 10 to incorporate a larger number of processors 12. Memory directories 18 within each processor 12 may support its associated local memory 16 and support a group of processors 12 within a local region depending on the desired size of each memory directory 18. Access to memory outside of a processor 12 or local region of processors 12 not supported by an individual memory directory 18 is handled by one or more external directories 22 and external switches 14. External switches 14 may also couple to input/output hosts 26 in order to support operations therewith. Each external switch 14 may also support processor network 28 extensions.

Thus, it is apparent that there has been provided in accordance with the present invention, a multi-processor system and method of accessing data therein that satisfies the advantages set forth above. Although the present invention has been described in detail, various changes, substitutions, and alterations may be readily ascertainable by those skilled in the art and may be made herein without departing from the spirit and scope of the present invention as defined by the following claims.